Using Lustre as a Wide Area Filesystem at 10Gb and Beyond

Stephen Simms ssimms@indiana.edu
Matt Link mrlink@indiana.edu
Robert Henschel henschel@indiana.edu
Michael Kluge michael.kluge@tu-dresden.de
http://pti.iu.edu/dc

N-Wave Stakeholder User Conference
May 10, 2011

INDIANA UNIVERSITY

The Data Capacitor Project

NSF Funded in 2005

535 Terabytes Lustre storage

24 Servers with 10Gb NICs

Short term storage

http://www.flickr.com/photos/shadowstorm/404158384/
http://www.flickr.com/photos/dvd5/163647219/
http://www.flickr.com/photos/vidiot/431357888/

Based on Lustre Filesystem

- Open Source
- Many thousands of client systems
- Petabytes of storage
- Over 240 GB/s I/O throughput at ORNL
- Scalable
  - aggregates separate servers for performance
- Standard POSIX filesystem interface

Data Capacitor as Central Filesystem

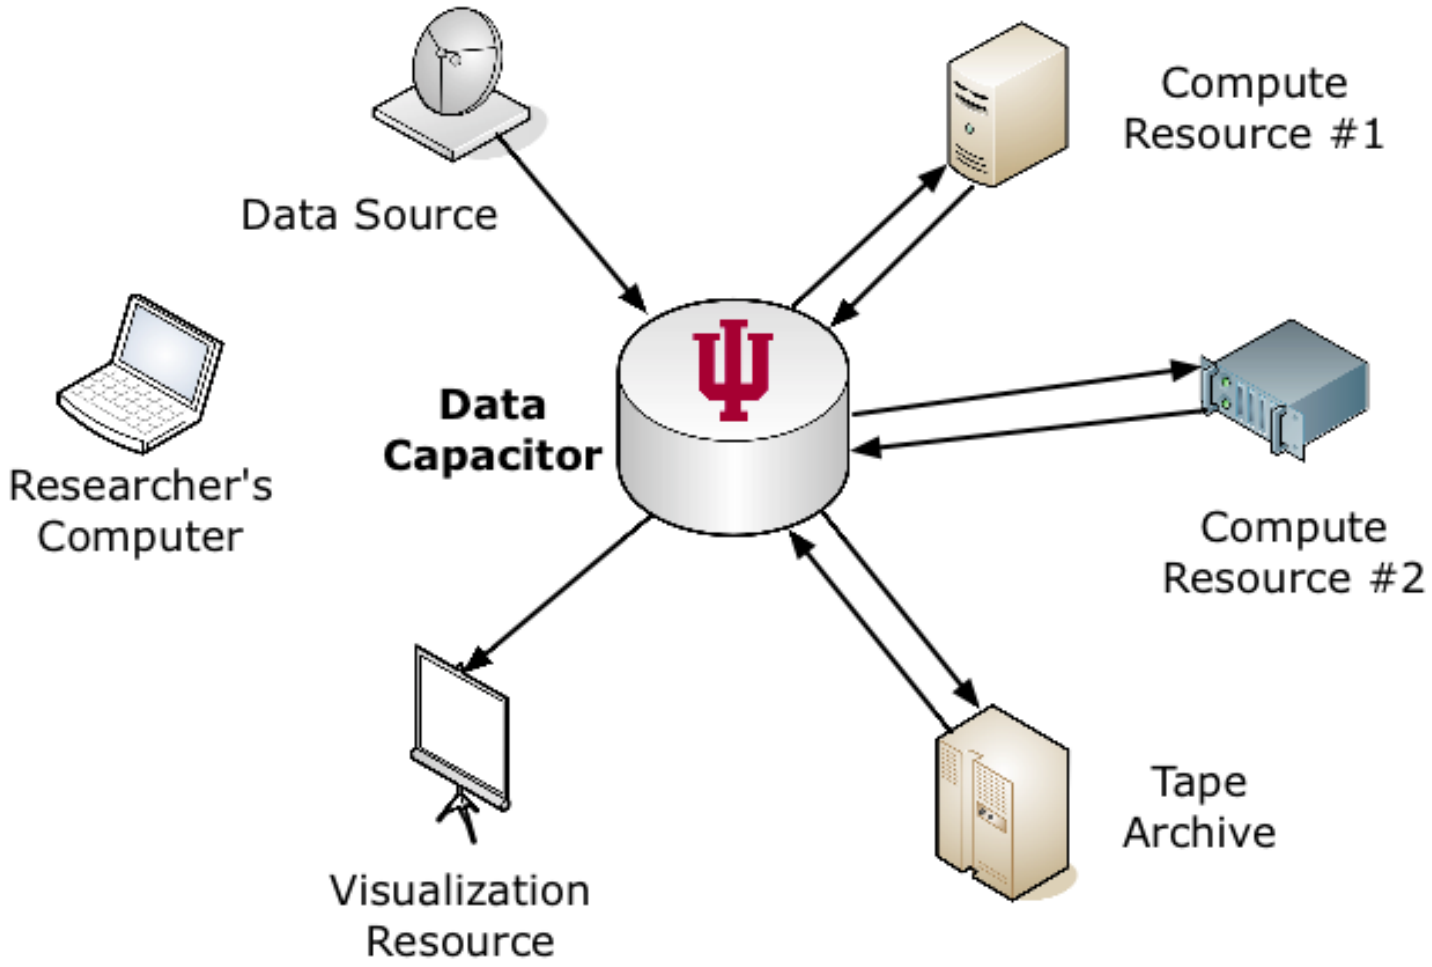

10 Gb Lustre WAN

977 MB/s between ORNL and IU
Using a single Dell 2950 client
Across 10Gb TeraGrid connection

2007 Bandwidth Challenge Win: Five Applications Simultaneously

- Acquisition and Visualization
  - Live Instrument Data
    - Chemistry
  - Rare Archival Material
    - Humanities
- Acquisition, Analysis, and Visualization
  - Trace Data
    - Computer Science
  - Simulation Data
    - Life Science
    - High Energy Physics

Bandwidth Challenge Configuration

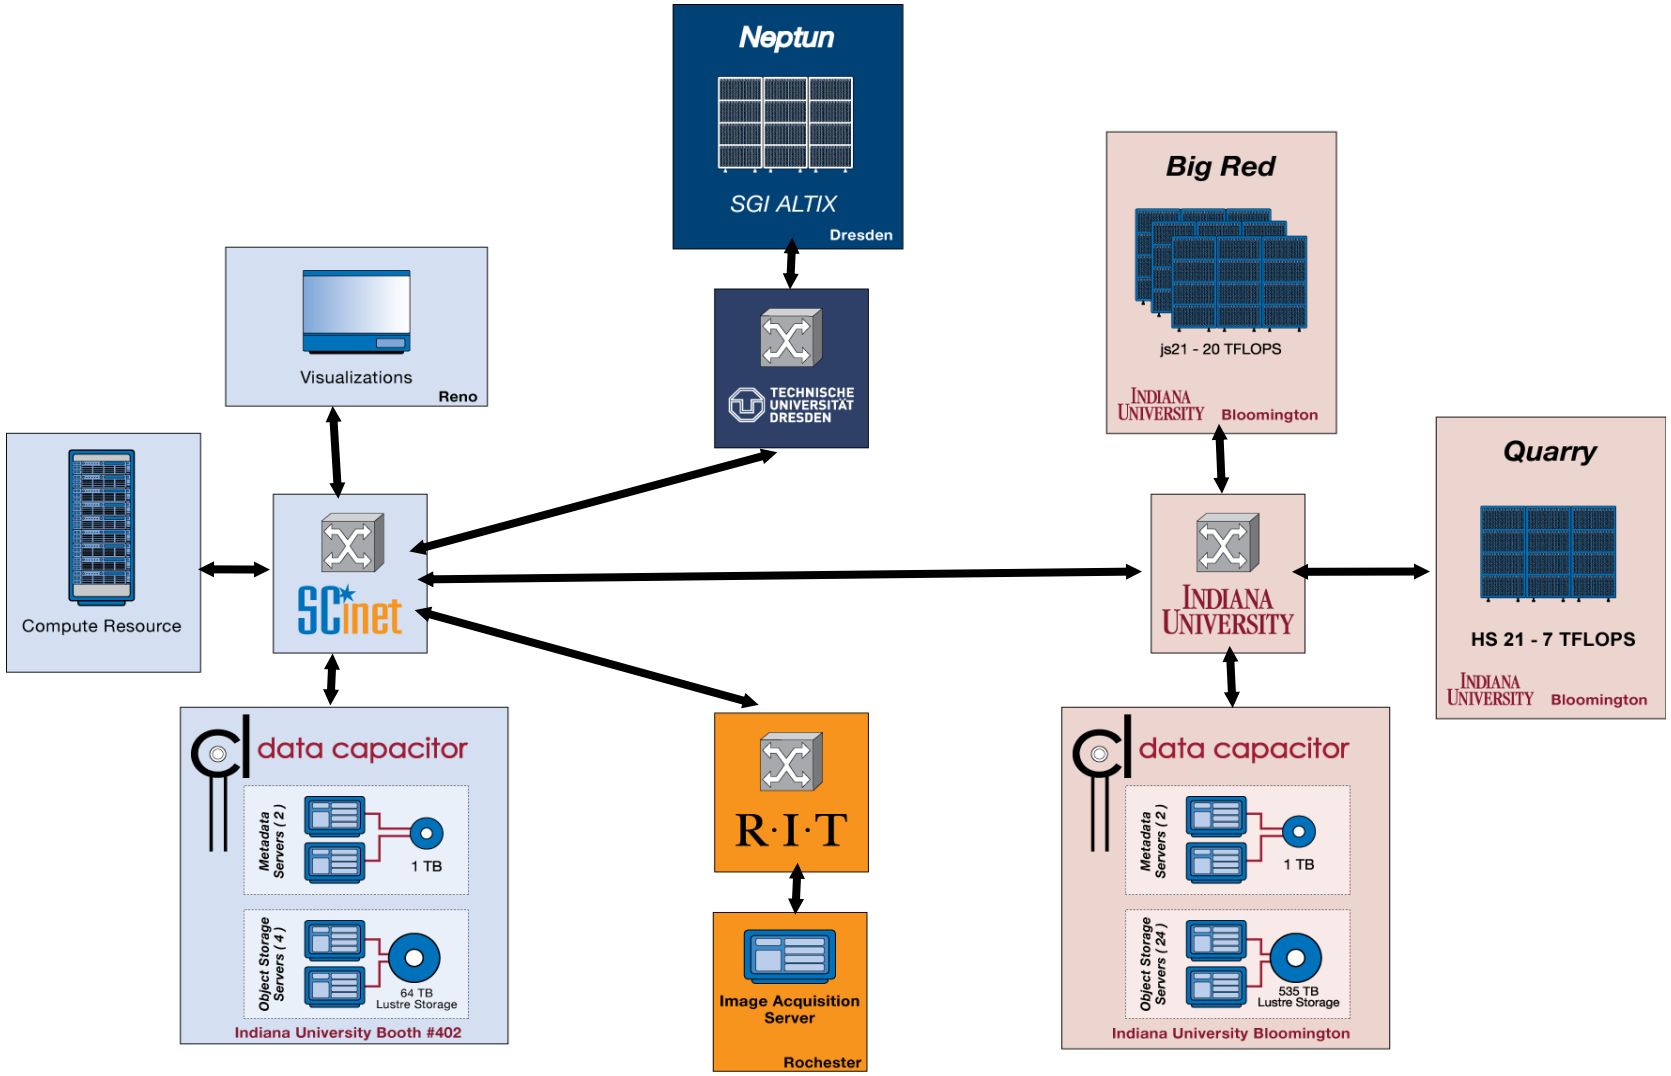

Challenge Results

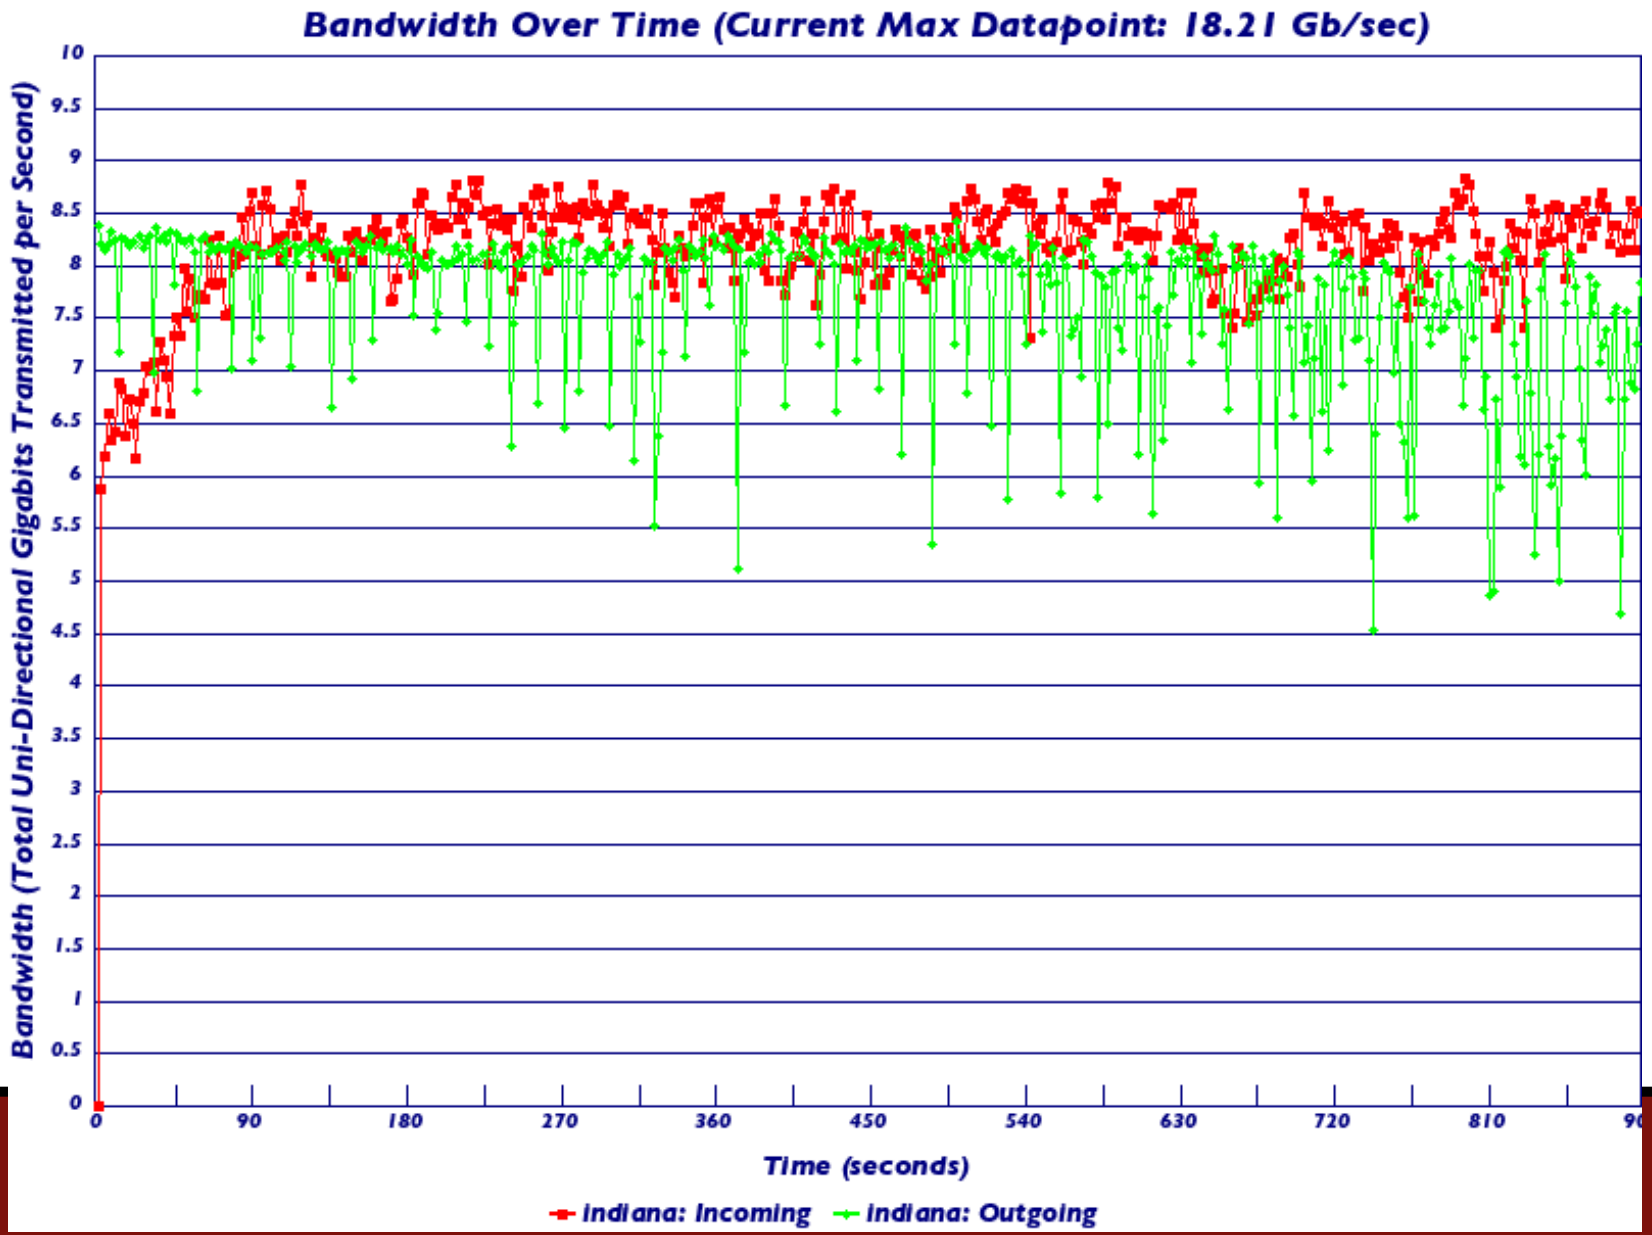

IU's Data Capacitor WAN Filesystem

- Funded by Indiana University in 2008
- Put into production in April of 2008
- 360TB of storage available as production service
- Centralized short-term storage for nationwide resources, including TeraGrid
  - Simplifies use of distributed resources
  - Projects space exists for mid-term storage

Transfer and Analysis of Ice Sheet Data (CReSIS)

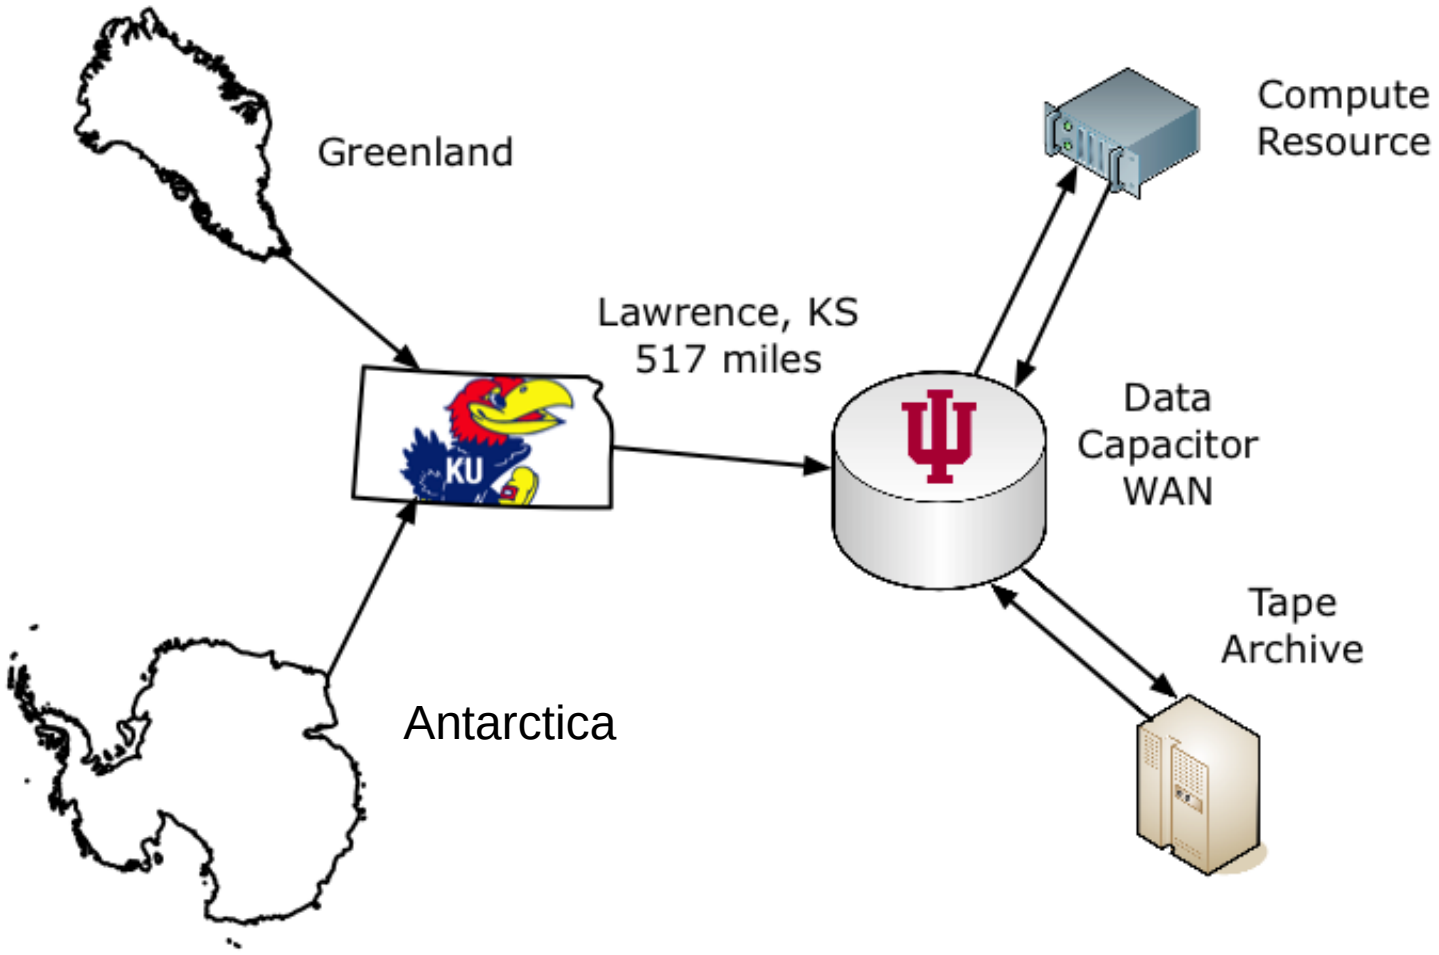

Dense Matter Research with DC-WAN

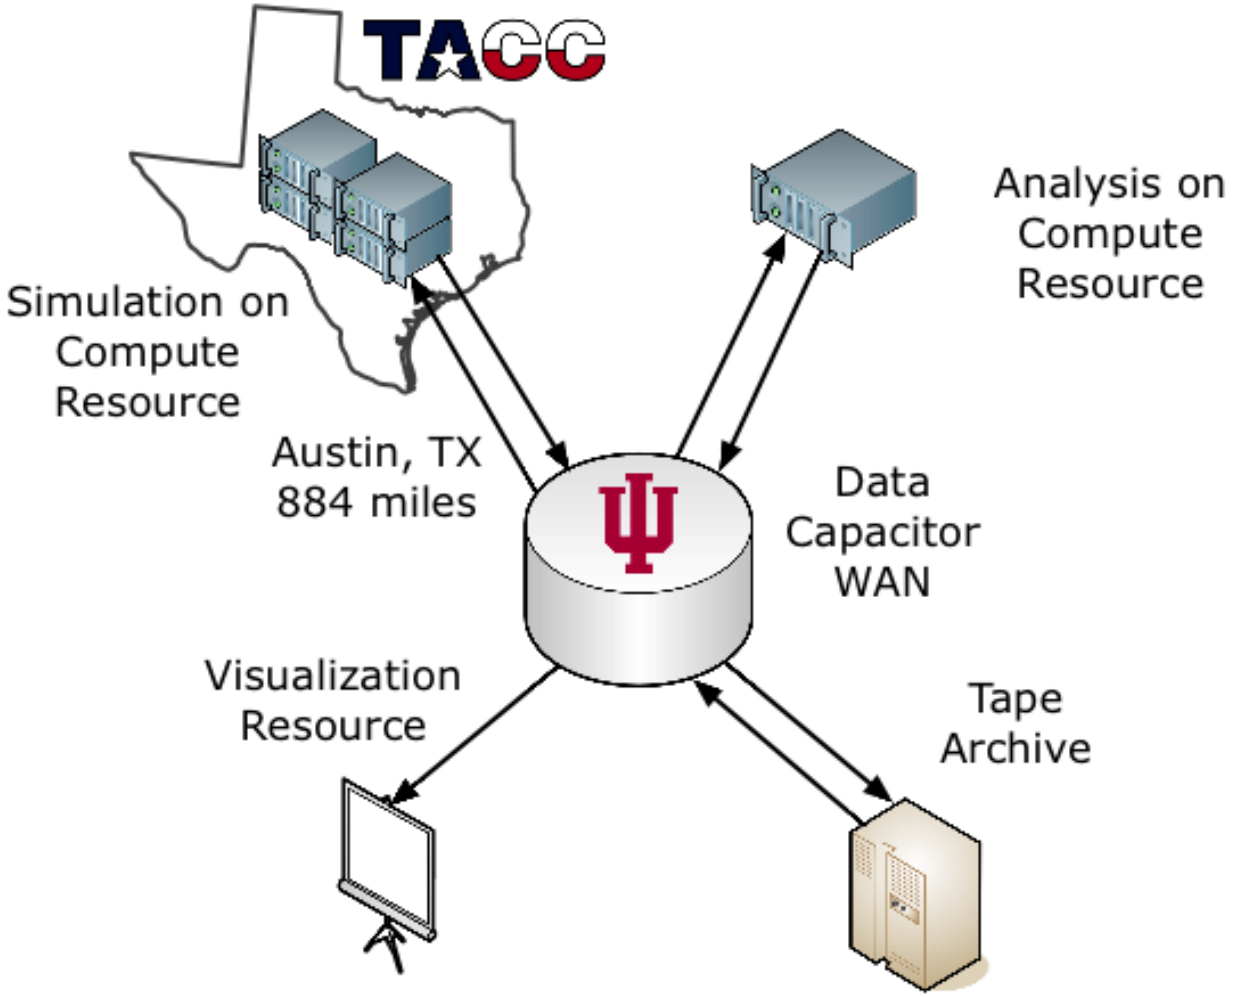

Gas Giant Planet Research

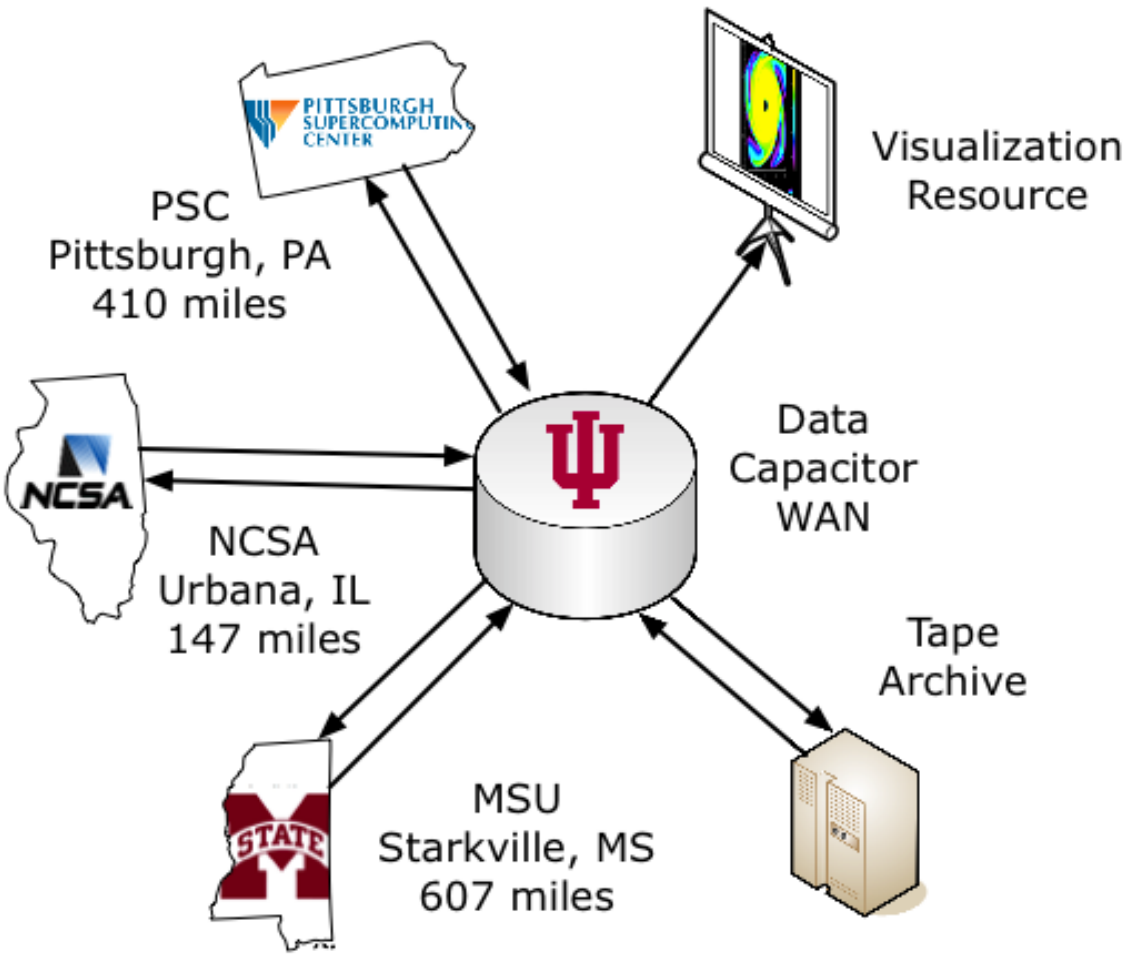

Data Capacitor WAN

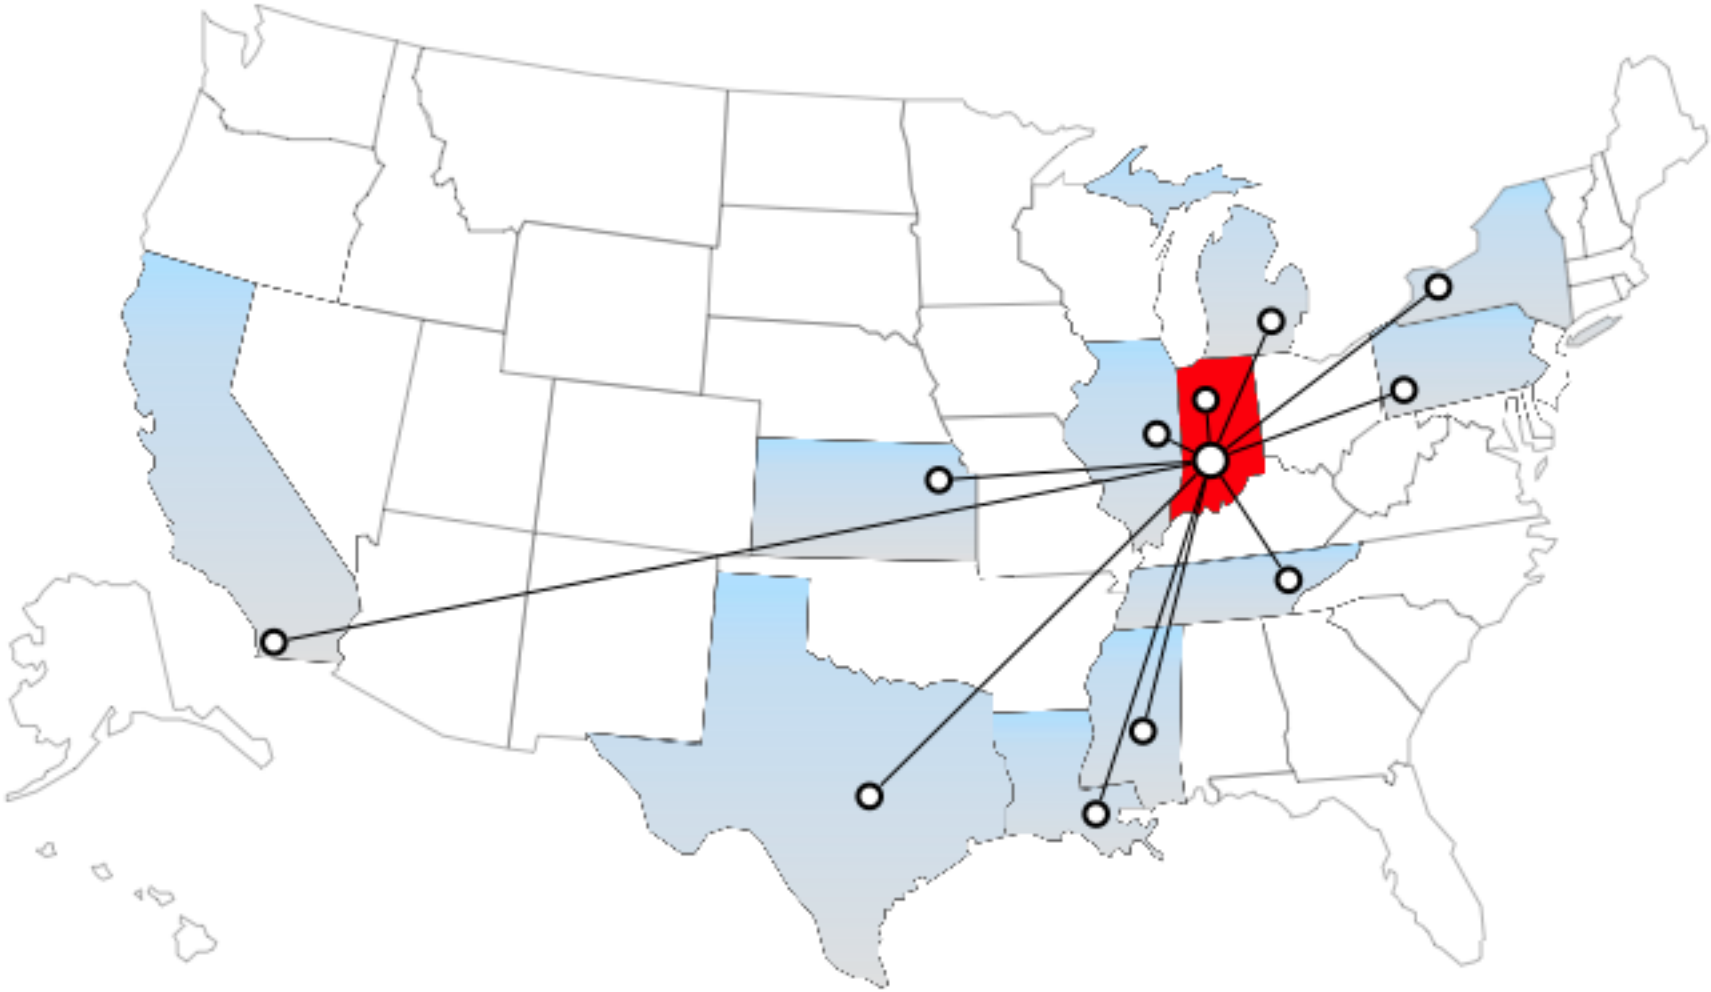

100 Gbit Testbed

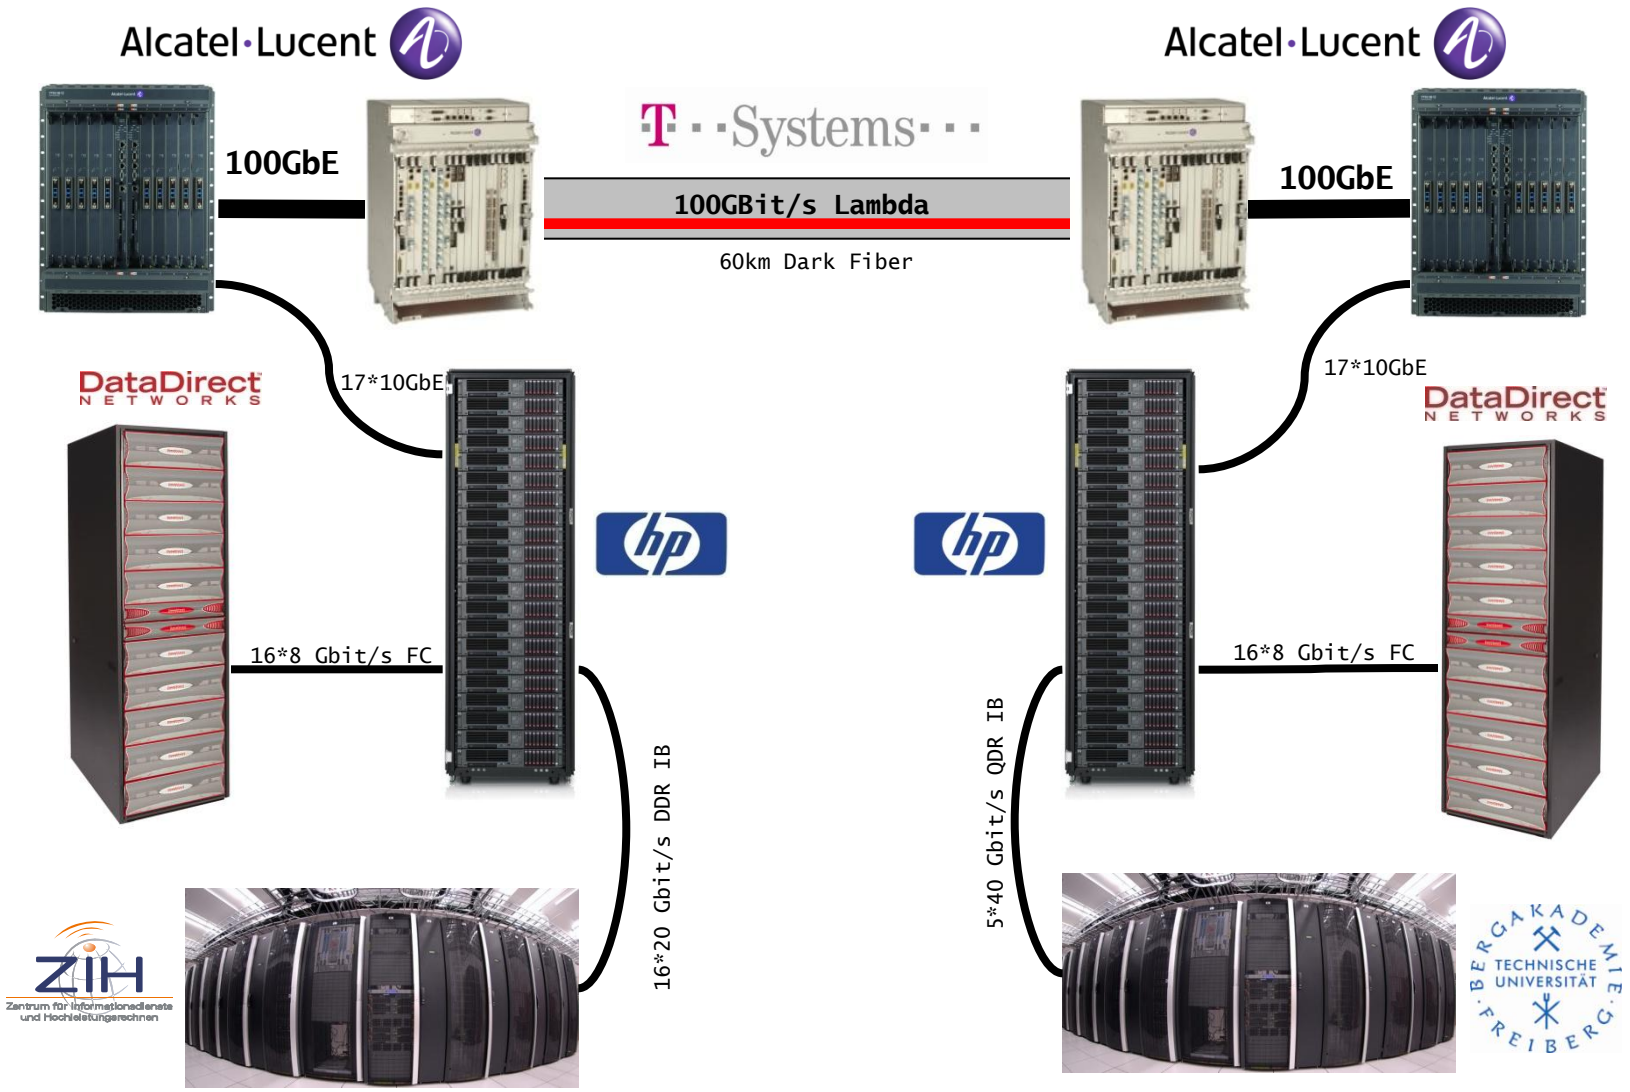

100 Gbit Testbed – Wide Area Project File Systems

100 Gbit Testbed – Uni-Directional Efficiency

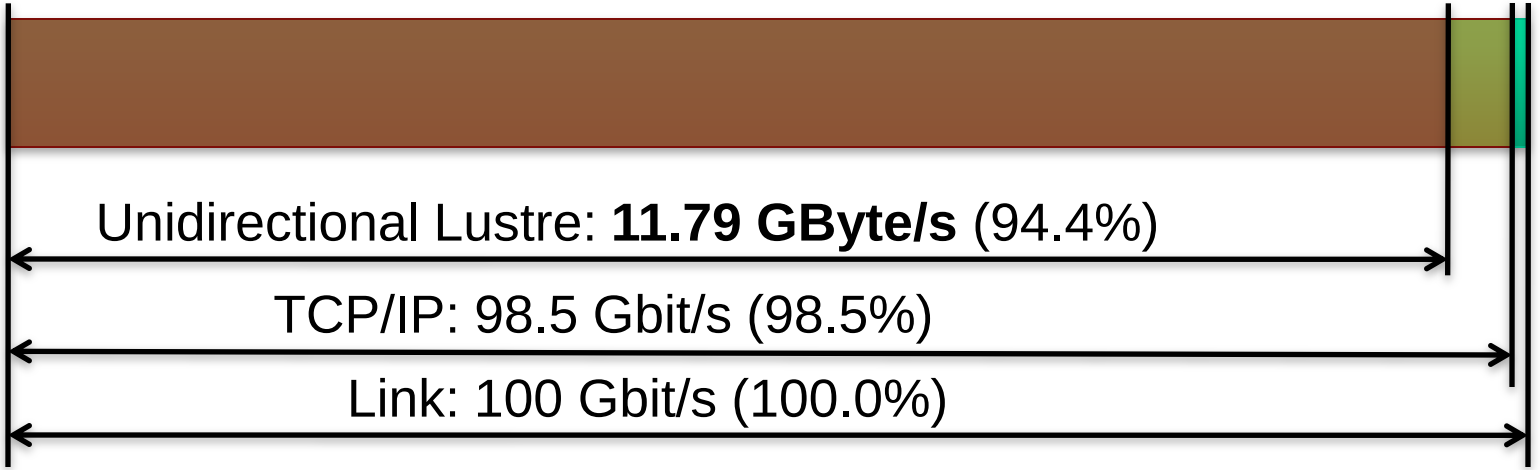

100 Gbit Testbed – Full Duplex Results

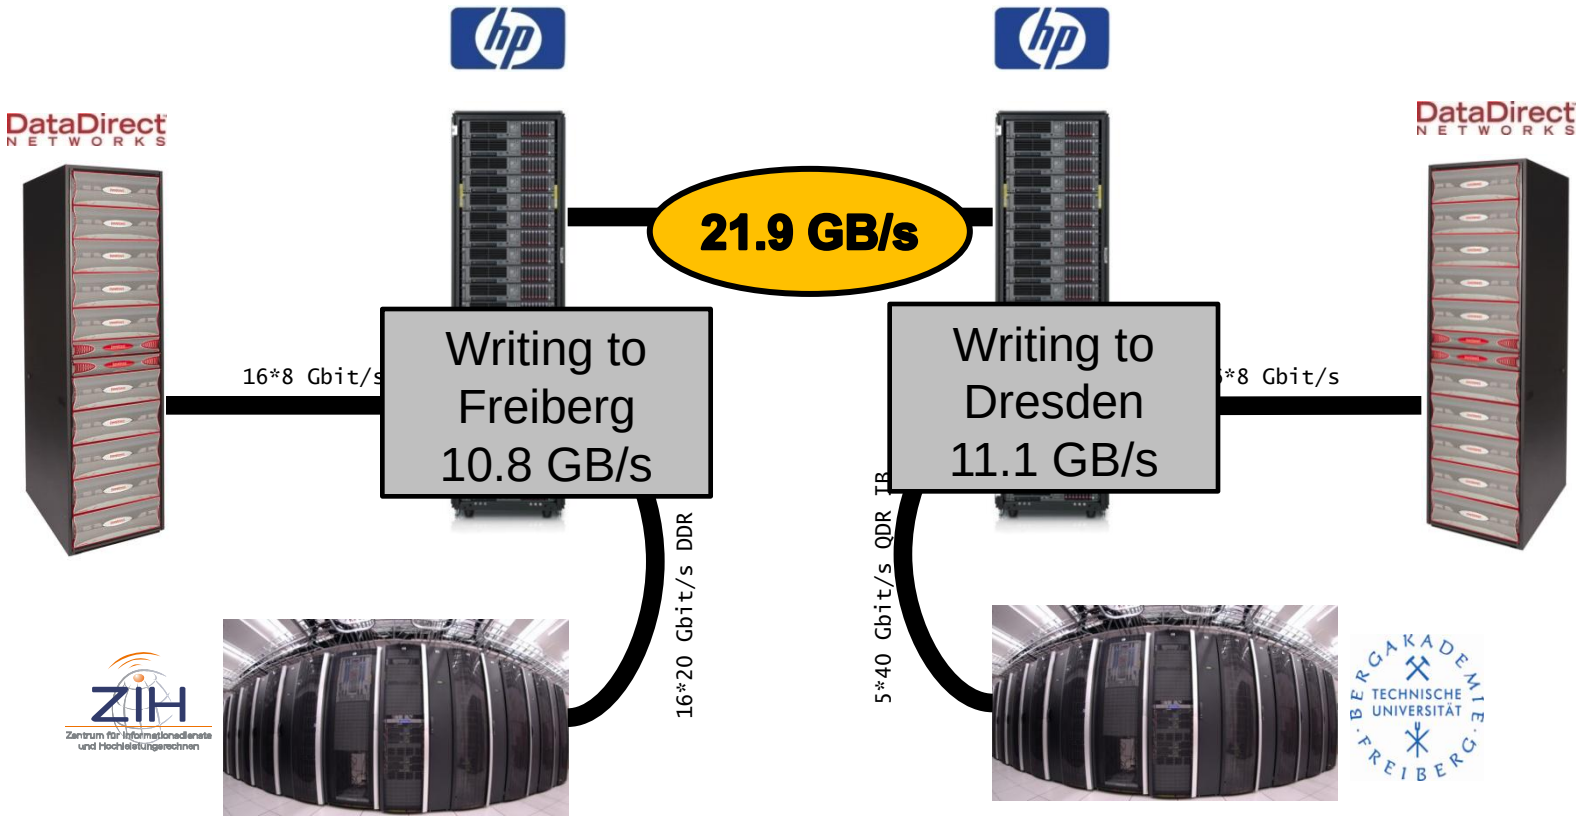

Acknowledgements

- Craig Stewart, Matt Davy, John Paul Herron, GRNOC (IU)
- IU's Data Capacitor Team
- IU's High Performance Systems Team
- The Research Teams and Sites Utilizing DC-WAN
- Data Direct Networks
- Eric Barton, Andreas Dilger, Brent Gorda (Whamcloud)
- Greg Pike (formerly ORNL)
- Doug Balog (formerly PSC)

This material is based upon work supported by the National Science Foundation under Grants No. CNS-0521433, ACI-0338618l, OCI-0451237, OCI-0535258, and OCI-0504075.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Thank you!

Stephen Simms
ssimms@indiana.edu

Team Data Capacitor
dc-team-l@indiana.edu
http://pti.iu.edu/dc